



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-TR-671579

Progress Report for "High-Resolution Methods for Phase Space Problems in Complex Geometry"

M. R. Dorr, J. A. F. Hittinger, P. Colella, P. W.
McCorquodale, P. O. Schwartz

May 28, 2015

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Progress Report for “High-Resolution Methods for Phase Space Problems in Complex Geometries”

M. R. Dorr (PI)¹, P. Colella², J. A. F. Hittinger¹, P. W. McCorquodale², and P. O. Schwartz²

¹Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, 7000 East Avenue L-561, Livermore, CA 94550.*

²Applied Numerical Algorithms Group, Lawrence Berkeley National Laboratory, One Cyclotron Road Mail Stop 50A-1148, Berkeley, CA 94720.[†]

May 28, 2015

1 Technical Progress

The goal of this project is to develop high-resolution numerical methods for the solution of continuum kinetic models. The motivating DOE application is the solution of gyrokinetic models of fusion reactor edge plasmas, which introduces unique geometric challenges. A key element of our approach has been the development, in tandem with our algorithmic research, of an experimental code named COGENT (Continuum Gyrokinetic Edge New Technology) that has been built on the Chombo adaptive mesh refinement library. In addition to providing a testbed for algorithm development, COGENT has served as a common platform for an ongoing collaboration with FES-funded fusion physicists under the banner of the Edge Simulation Laboratory (ESL) (see Section 4). In the following subsections, we describe our progress on multiple fronts within the past year.

1.1 Publication of high-order, mapped-multiblock, finite-volume discretization paper

Our paper [6] describing the extension to multiblock grids of our high-order, mapped-grid, finite-volume discretization approach [4] was accepted by the Journal of Computational Physics. The key idea of the paper is the construction and use of sufficiently high-order interpolation to fill ghost cells at the interfaces between blocks, followed by the use of the interior scheme of [4] all the way up to the block interfaces. Three example problems are included that demonstrate the effectiveness of the approach, including one motivated by the edge plasma problem. Although the referee reports did not request that we re-run or extend any of our computational results, they did request additional information in the revision, which consumed a substantial amount of time to assemble. This request included a more extensive comparison with prior mapped grid discretization approaches that have been used extensively in the aerodynamics community. Additional text was added to clarify the range of applications for which the approach may be appropriate. For example, although we have, in fact, used the discretization approach described in this paper for elliptic problems, we elected to limit the scope of the current article to hyperbolic conservation laws to avoid further delay in publication.

*This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

[†]Research supported by the Office of Advanced Scientific Computing Research of the US Department of Energy under contract number DE-AC02-05CH11231.

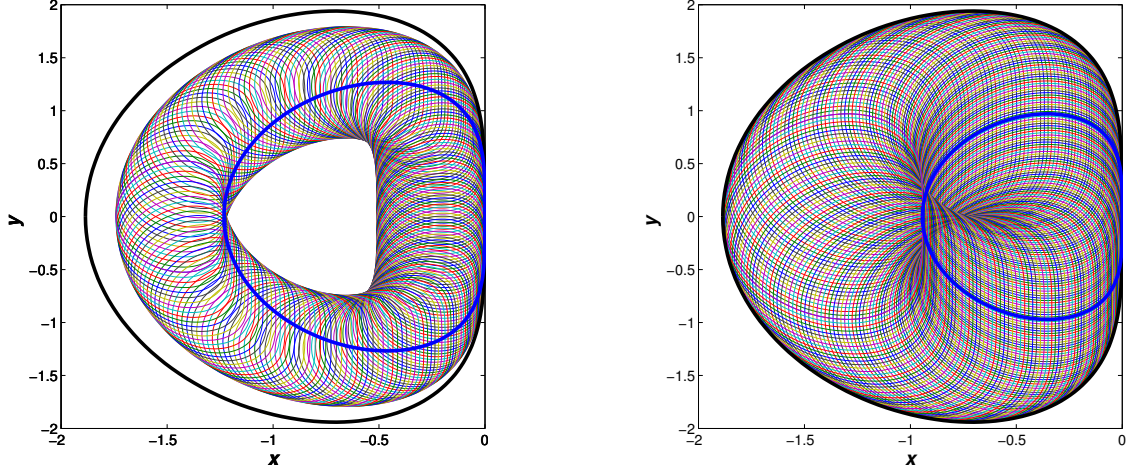


Figure 1: Envelopes of the eigenvalues for the third-order upwind scheme for $D = 2$ in the complex plane for two different flow angles: $a_1 = |a| \cos(\pi/8)$ (left) and $a_1 = |a| \cos(\pi/4)$ (right). The curve $a_1|a|^{-1}g_3(\phi)$ is the thick blue line in each case, and the bounding envelope for all flow angles, $\sqrt{2}g_3(\phi)$, is the thick black line.

1.2 Stability analysis of multi-dimensional time integration algorithms

Input from our physics collaborators indicated that, for many problems, they could take a larger time step than that dictated by our original multidimensional stability estimates. In addition, in the process of debugging, we had implemented first-, third-, and fifth-order linear upwind fluxes, as well as the fifth-order WENO flux, but had not revisited the time step criteria. Each of these schemes has different time step stability requirements, so understanding the algorithm behavior required us to revisit the time step stability conditions.

For a multidimensional linear advection equation with velocity $a = (a_1, a_2, \dots, a_D)$, an upwind discretization of the spatial operator has eigenvalues

$$\lambda_j = -\frac{|a|}{h} \sum_{d=1}^D \frac{a_d}{|a|} g_p(\theta_{j_d}), \quad j_d \in [1, N] \quad (1)$$

where h is the mesh size (with N the number of cells in each direction), $\theta_{j_d} = 2\pi(j_d - 1)/N$, and $g_p(\theta)$ corresponds to a complex function specific to the upwind scheme of order p . A simple trick used in the analysis is that $|a_d/a| \leq 1$ for all d , so $|\lambda_j| \leq h^{-1}D|ag_p(\phi)|$, where ϕ is the continuous extension of θ_{j_d} . Unfortunately, this gives only a loose bound.

Revisiting this issue, the goal was to determine the envelope of the θ_j and then to find the scaling of the eigenvalues that ensures that this envelope lies within the stability region of the time integration scheme. We realized that, for each d' , the $g_p(\theta_{j_d})$ for $d < d'$ merely describes a vector in the complex plane by which the origin of g_p is displaced; for d' , $g_p(\phi)$ traces the locus of eigenvalues starting from this shifted origin. Consider the case of $D = 2$: the locus lies on the boundary of the complex domain where, at each point on the locus for $d = 1$ (scaled by the direction cosine), the (scaled) locus for $d = 2$ is traced. Examples are shown in Figure 1. Using the Envelope Theorem [5], one can prove that this combination of the two loci is bounded by the same function, appropriately scaled. The method of Lagrange multipliers and the constraint that $\sum_d a_d^2 = 1$ produces a tighter bound that occurs when each dimension is scaled equally, i.e., $|a_d/a| = \sqrt{D}/D$, which is intuitively advection along the multidimensional diagonal. The extension to $D > 2$ can be made by induction. Thus, a tighter bound on the multidimensional eigenvalues is $|\lambda_j| \leq h^{-1}\sqrt{D}|ag_p(\phi)|$, which is generic to any upwind scheme.

We have successfully implemented time step criteria based on this analysis. To the best of our knowledge, this is a new theoretical result, broadly applicable to any multidimensional, linear upwind discretization

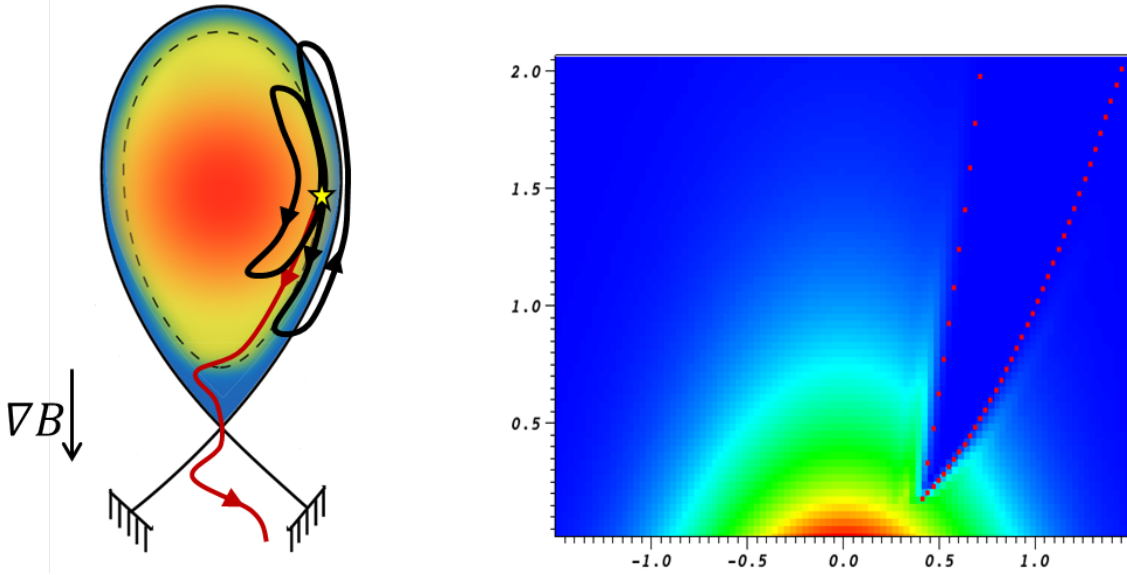


Figure 2: Schematic of loss cone test problem (left) and distribution function velocity variation (right) at the point indicated by the yellow star. The red dots mark the loss cone boundary predicted by an analytic model.

scheme. We have heard of instances where a similar criteria has been used in practice, but have yet been able to find or be directed to a rigorous proof. We are currently preparing a manuscript for publication on this work.

1.3 Algorithm verification on a multiblock plasma edge geometry

Using our testbed code COGENT, we continued to conduct verification studies of our high-order, mapped-multiblock, finite-volume discretizations and solution algorithms on the edge plasma problems that initially motivated this research. One of our most important ongoing test cases is the so-called “loss cone” problem. Although we have previously reported qualitatively correct results for this problem, within the past year, in collaboration with our FES collaborators, we have now demonstrated quantitative agreement with an analytic model in an analytically prescribed geometry.

Without going into all of the details, the basic behavior of this problem is that plasma species particles trapped by the magnetic field follow banana-shaped orbits, as indicated by the cartoon in the left-hand side of Figure 2. At points near the location of the yellow star, the orbits span the magnetic separatrix between the closed and open field line regions. Depending upon the particle energy, particles drifting into the open field line region are either “scraped-off” toward the divertor plates (denoted by the line segments at the bottom of Figure 2 (left) or remained trapped, traveling back into the closed field line region that continues to be replenished from the core (via an applied boundary condition in COGENT). The net result is the voiding of a velocity space “cone”, as shown in the plot in the right-hand side of Figure 2, which shows the velocity variation of the distribution function at the configuration space point indicated by the yellow star. The red dots show the loss cone boundary predicted by the analytic model, which agrees quite well with the COGENT solution.

1.4 Discretization near the X-point

We have encountered a challenging question with respect to the use of discretization in a mapped coordinate system to address solution anisotropy. In the edge plasma application, the motivation for using a coordinate mapping aligned with magnetic field lines is the hope that a corresponding grid can be deployed whose cell

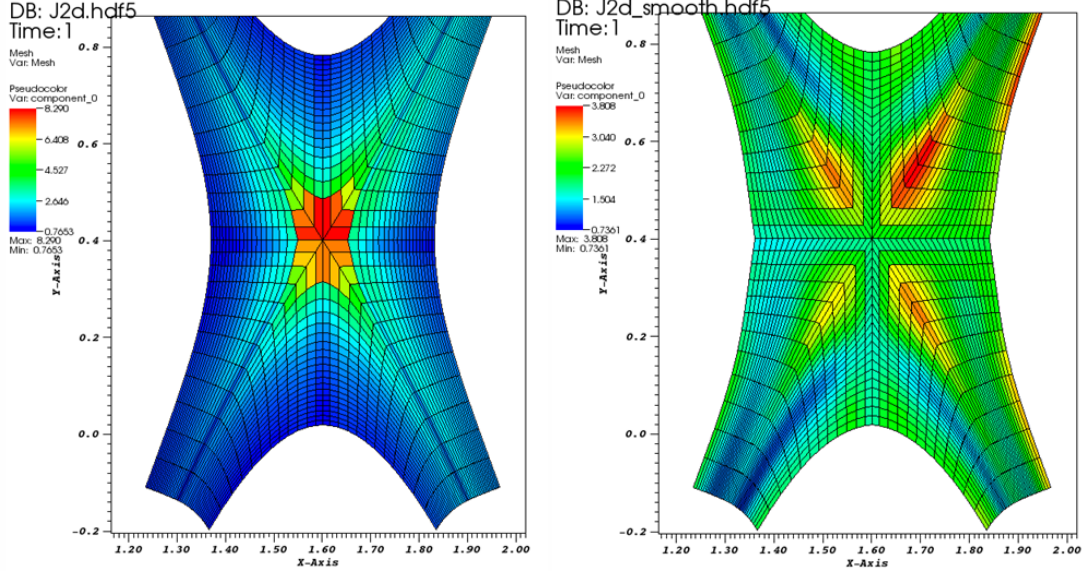


Figure 3: Comparison of mapping Jacobian on field-aligned but less smooth (left) versus less field-aligned but smoother grids (right). Note the difference in scales from left to right is roughly a factor of two.

widths along field lines can be substantially larger than in the perpendicular direction, thereby reducing the total number of cells. However, as we have noted before, the use of field-aligned mappings all of the way up to the X-point is problematic due to the increasingly large derivatives appearing in the metric factors. This also precludes the smooth extension of block mappings through the X-point, as required by our high-order, mapped-multiblock, finite-volume discretization [6].

We are therefore compelled to abandon field alignment in a neighborhood of the X-point, the extent of which then becomes a potentially important algorithmic parameter. As shown in the plots in Figure 3, depending upon the size of the neighborhood, the resulting grids range from those that are more field-aligned but less smooth (left) to those that are less field-aligned, but smoother (right). The colors in the Figure 3 plots indicate the magnitude of the associated Jacobian mapping. Unlike the left-hand plot, in which the largest Jacobian magnitude occurs exactly at the X-point, the mesh in the right-hand plot has a much smaller Jacobian at block boundaries. Because of this, the grid extensions resulting from the left-hand mapping are quite poor compared to those resulting from the right-hand mapping, which yields significantly better behavior of the discretization algorithm.

One might look to the use of some physical criteria to help guide the choice of the X-point dealignment neighborhood size. For example, since we are concerned about flow anisotropy, the ratio v_{\perp}/v_{\parallel} of perpendicular to parallel velocity is an appropriate measure. However, for typical tokamak parameters, this criteria suggests the use of a very small dealignment neighborhood, resulting in mappings/grids that look more like the left-hand plot of Figure 3 than the more desirable right-hand plot. In other words, solution anisotropy remains a concern much closer to the X-point than we had expected.

As a result of this analysis, we are investigating options to maintain field alignment closer to the X-point. We have performed several studies demonstrating that, if field alignment is abandoned too far from the X-point, finer gridding is required to resolve the solution anisotropy that is not efficiently represented. Although this might seem somewhat obvious, the important fact is that we now have a much better quantitative sense of how these numerical issues interact. A key observation is that, whereas the metric factors of a field-aligned mapping become large near the X-point, the various integrals (*e.g.*, over control volume faces) appearing in a finite-volume discretization always remain finite. Moreover, the distribution function admits no singular behavior near the X-point. We have therefore begun to investigate the use of alternative means to integrate the metric factors near the X-point accurately, accounting for their near singularity on more field-aligned grids. Because the poloidal component of the magnetic field vanishes at the X-point, one approach we will

pursue is to use a Taylor series expansion in a neighborhood of the X-point to provide analytic expressions for use in the metric factor integrations.

1.5 Solution of the gyrokinetic Poisson-Boltzmann system

A major source of fast timescale concerns in the edge plasma application comes from the inclusion of electrons, which stream rapidly along magnetic field lines. Because the details of the electron dynamics are often relatively unimportant, we have been pursuing the use of a Boltzmann electron model in the gyrokinetic Poisson system:

$$\nabla \cdot \left\{ \left[\epsilon_0 \mathbf{I} + e^2 \sum_{\alpha} \frac{Z_{\alpha}^2 \bar{n}_{\alpha}}{m_{\alpha} \Omega_{\alpha}^2} (\mathbf{I} - \mathbf{b} \mathbf{b}^T) \right] \nabla \Phi \right\} = e \left(\frac{\langle \sum_{\alpha} Z_{\alpha} \bar{n}_{\alpha} \rangle}{\langle \exp(e\Phi/T_e) \rangle} \exp(e\Phi/T_e) - \sum_{\alpha} Z_{\alpha} \bar{n}_{\alpha} \right). \quad (2)$$

The $\langle \cdot \rangle$ in the Boltzmann relation prefactor (coefficient of the first term in the right-hand side) denotes averaging over magnetic flux surfaces. This electron model assumes that electrons instantaneously redistribute themselves to maintain charge neutrality on the flux surfaces. This form of the prefactor is only valid on closed field lines in the plasma core; a different form is needed to describe the open field line region outside of the magnetic separatrix.

Although the use of a Boltzmann electron model eliminates the need to resolve a fast time scale, as with many models constructed to step over stiff processes, the numerical challenge then shifts to the solution of the resulting nonlinear system. Our initial attempt to solve (2) utilized a standard Newton iteration with a multigrid-preconditioned Jacobian solver. Reflecting the natural decomposition of the problem into parallel (*i.e.*, in the field direction \mathbf{b}) and perpendicular components, a preconditioner was found that worked well in annular geometries where the solution variation along field lines was small. Unfortunately, in other cases, the solution obtained was unsatisfactory, in spite of the fact that the Newton solver seemed to have converged. We determined the cause of this behavior to be the relative scaling of the terms of (2). In the core region (of a tokamak like the DIII-D machine at General Atomics) the left-hand side of (2) is three orders of magnitude smaller than the right-hand side. In the open field line region, the difference is two orders of magnitude. Moreover, if one considers the asymptotics of the variation of the potential Φ relative to its average value on flux surfaces, one discovers that in the core region the potential variation is much smaller than its average. The net result of these observations was that (2) cannot be solved using a vanilla Newton approach. We therefore developed a solution strategy that accounts for these important scaling differences, iterating a parallel solve (*i.e.*, a fixed point iteration within field lines) and a perpendicular solve (*i.e.*, a tridiagonal system across field lines) to self-consistency.

1.6 Consistent magnetic geometry

We found a solution to the problem of constructing a smooth, divergence-free magnetic field that is consistent with the coordinate mappings defining the COGENT mesh. Although this might seem like a somewhat mundane technical detail, since a major focus of this project is the use of mapped coordinates to address magnetic-field-induced anisotropic flow, the careful specification of the field and block mappings is, in fact, quite important. For our fourth-order, finite-volume discretization of the gyrokinetic system, we require a smooth (at least five continuous derivatives) representation of the magnetic field. Moreover, because we insist on a conservative formulation, we require that the gyrokinetic phase space advection velocity be divergence free, at least to fourth-order accuracy. Since a divergence-free magnetic field implies a divergence-free velocity, we desire a discrete magnetic field whose divergence is as small as possible.

Because COGENT uses magnetic field data obtained by interpolation from actual experimental data, we had previously encountered unacceptably large errors introduced by large magnetic field divergences. We therefore employed a projection method to subtract out the undesirable gradient components. This approach was not entirely satisfactory for two reasons. First, there was no guarantee that projection-corrected field was still well aligned with the coordinate mapping that had already been generated using the original field. Secondly, the field derivatives (*e.g.*, for the field curl and magnitude gradient) could not be computed to the desired fourth-order for reasons that are too detailed to discuss here. We briefly experimented with the use of

divergence-free radial basis function interpolations as well as trigonometric interpolations obtained from FFT coefficients of the interpolated field data, in which the gradient component was removed in Fourier space. Although the latter approach addressed the divergence-free issue, there was still the matter of maintaining consistency of the resulting field with the block coordinate mappings.

As a consequence of our involvement with the ATOM SciDAC FES Partnership project (see Section 4), we learned of a code named Hypnotoad [1], which is being developed by Ben Dudson at the University of York and his collaborators for use in the BOUT++ plasma fluid code. Hypnotoad is a field-aligned mesh generator for tokamaks, which utilizes experimental field data in a particular standardized format. A main component is the calculation of level surfaces of the magnetic flux ψ , which Hypnotoad constructs as a discrete cosine transform (DCT) expansion. In addition to using the field-aligned mesh produced by Hypnotoad to generate our block mappings via high-order interpolation, we made a minor modification to Hypnotoad to extract the DCT coefficients for use in reconstructing ψ in both COGENT, as well as the mesh smoothing code discussed in Section 1.7. Not only is the flux ψ constructed in this way analytic, but the corresponding axisymmetric field $\mathbf{B}(R, Z) = (-\partial\psi(R, Z)/\partial Z, \partial\psi(R, Z)/\partial R)/R$ is patently divergence-free.

1.7 Coordinate mapping generation

We have continued our development of algorithms to generate coordinate mappings for use in our mapped-multiblock, finite-volume discretizations. In addition to requiring that mappings possess sufficient smoothness (*e.g.*, at least four continuous derivatives for a fourth-order scheme), other application-specific properties may also need to be enforced. In the edge plasma application, we require flux surface alignment, *i.e.*, flux surfaces are parameterized by one of the mapped coordinates. We are therefore pursuing an optimization-based strategy that minimizes a functional of the coordinate mappings in which deviation from smoothness and flux surface alignment are penalized. The critical point of the optimization problem is obtained by solving an Euler-Lagrange equation on an appropriate grid. A high-order (B-spline or spectral) interpolant of the resulting solution then provides the coordinate mapping.

In the past year, we resolved a number of technical challenges that arose when pursuing this strategy. In our initial formulation of the algorithm, we imposed the smoothness penalty (essentially a minimum curvature condition) in all coordinate directions simultaneously. Since the flux surface alignment penalty term acts only in one coordinate direction, this approach introduced an unstable coupling in the system. This was rectified by treating the objective function contributions from each coordinate direction separately.

Another complication resulted from the fact the residual of the Euler-Lagrange system obtained by this approach consists of potentially large terms (resulting from the smoothness and alignment terms in the objective function) of opposite sign. Depending upon the choice of objective function weights, it is therefore possible to conclude that the Euler-Lagrange residual is small, when in fact the individual constraints are far from satisfied. We have developed a strategy to detect and mitigate such situations, but more testing is required as well as further refinement to reduce the number of free parameters being set.

To utilize the mapped multiblock algorithm described in [6], we require smooth extensions of the block mappings beyond block boundaries. This places a further constraint on the mapping generation problem. As a first step, we have developed smooth extrapolations of the block mappings obtained by the optimization approach described above. However, we would ideally like to include the (smoothness of) the extrapolated mappings as part of the optimization process itself. More generally, we hope to focus our optimization-based approach to the mapping generation problem more directly at the specific requirements of the mapped-multiblock discretization, rather than a general smoothness criteria.

2 Noteworthy changes in scope, schedule, budget, issues, or personnel in the last year or anticipated in future

Our renewal proposal included as a deliverable the investigation of implicit-explicit time integration methods to address the multiple time scales arising in the kinetic models that we have been investigating. Although the stage had already been set with respect to the implementation (by Jeff Hittinger) of an IMEX framework

within COGENT in the previous funding cycle, we have not as yet made significant further progress on this topic due to a manpower shortfall and some re-prioritization. In addition to Jeff’s other responsibilities, including his serving as the PI of the ExReDi project, his work described in Section 1.2 on the development of the multidimensional stability analysis for the explicit time integration algorithm consumed most of the time he had available to work on this project. We felt that it was more important to nail down this very important aspect of the explicit algorithm before continuing the IMEX investigation. Moreover, given the applicability of the new stability result to multidimensional problems well beyond the original scope of this project, acceleration of the publication of this work seemed well justified.

We also conducted a search for an appropriate postdoctoral appointee to assist in the investigation of time-integration and high-order discretization algorithms, and we have been very fortunate to have found an extremely well-qualified candidate. Dr. Debojyoti (“Debo”) Ghosh is currently a postdoc at Argonne National Laboratory, who has accepted an offer to join our project this coming September. Debo has been working with Emil Constantinescu and Jed Brown on IMEX integration schemes and high-order compact essentially non-oscillatory schemes for compressible and incompressible atmospheric flows. Although we would welcome Debo’s arrival immediately, he requested a September start in order to complete his current work at Argonne. Debo does not have any experience with our computational plasma physics application, but we are confident he will come up to speed quickly. His expertise in time-integration and high-order spatial discretization fit our current needs very well, and should greatly accelerate our progress.

3 Workplan for remainder of FY15 and beyond

In addition to the ongoing work discussed in the Progress section above, when our new postdoc arrives in the Fall, we will resume our investigation of IMEX schemes by continuing to study the implicit integration of collision operators using an additive Runge-Kutta (ARK) strategy. Part of the rationale behind considering collisions first is the variety of collision operators available in COGENT (ranging from simple Krook operators to fully nonlinear Fokker-Planck operators) and the relative simple velocity space geometry (*i.e.*, a Cartesian grid). The verification of stability at the desired time steps and the preservation of order of the combined operators will continue into FY16. We will also begin to analyze the effectiveness of ARK methods to address the additional fast time scales resulting from reduced electron models. Specifically, we will investigate the performance in the ARK setting of the edge plasma-relevant Boltzmann model proposed by our ESL physics collaborators in [3] and possibly an inertia-less fluid model as well.

We will also continue our work on the development of a nonlinear error transport model. Because of the method-of-lines formulation of our high-order schemes, the construction of the error transport operator can be obtained through differencing the primal transport operator applied to the corrected primal solution and applied to the uncorrected primal solution [2]. Incorporating this into COGENT will be relatively straightforward. The more challenging application of the error transport technique will be in the discrete evaluation of the residual of the primal solution, which must be done to at least one order higher than the original primal discretization. There are several approaches to this, and we will consider both a higher-order least-squares reconstruction on the physical domain as well as a computational-grid formulation that will require sixth-order finite volume operators and mapped grid metrics. Experience with nonlinear error transport has shown that higher-order boundary conditions, while not an insurmountable obstacle, will probably require some substantial effort to implement. We will verify the implementation of the error transport capability by using manufactured solutions.

4 Collaborations and relationships with other SC offices / projects

This project is a collaboration between members of the Center for Applied Scientific Computing (CASC) at LLNL (Dorr and Hittinger) and the Applied Numerical Algorithms Group (ANAG) at LBNL (Colella, McCorquodale and Schwartz). All of the investigators are actively involved in algorithm development and collaborate closely on software implementation. The LLNL investigators have been the primary developers

of COGENT. As the developers of the Chombo adaptive mesh refinement library, upon which COGENT completely depends, the LBNL investigators have the primary responsibility for extending the Chombo infrastructure to support the specialized needs of this project not addressed by the ongoing development of Chombo under FASTMath for a broader range of applications.

This numerical methods research project is part of a multidisciplinary collaboration with the DOE Office of Fusion Energy Sciences (FES) theory program called the Edge Simulation Laboratory (ESL). The FES component funds plasma physicists at LLNL (Mikhail Dorf, Thomas Rognlien, Andris Dimits), General Atomics (Phillip Snyder, Jeffrey Candy, Emily Belli) and the University of California at San Diego (Sergei Krasheninnikov’s students) to pursue various aspects of edge plasma science. Partly through the joint development of COGENT, our ESL collaborators have therefore been the first adopters of the methods and software developed by this ASCR project. Conversely, they have helped define appropriate test problems and verify the COGENT implementation.

This project also leverages work performed by the PI as part of the FASTMath SciDAC institute. One of the goals of FASTMath is to improve the interoperability of SC-developed software. To apply the high-order, mapped-multiblock, finite-volume discretizations being developed by this project to problems employing implicit time integration or with constraints, linear operators (matrices) must be constructed in order to call solvers from packages such as *hypre* or PETSc. In FASTMath, we are developing an interface between the Chombo mapped-multiblock data objects and both *hypre* and PETSc. In this way, the research results of this base program research will be made more accessible to a broader community.

In late FY14, a new SciDAC partnership project named AToM (Advanced Tokamak Modeling) was jointly funded by Fusion Energy Sciences and ASCR. The goal of AToM is to enhance the predictive modeling capabilities that currently exist within the US magnetic fusion program through the expanded development of a production workflow manager and high-performance computing framework. In the first six months of the project, work has begun to integrate a variety of codes under the One Modeling Framework for Integrated Tasks (OMFIT) workflow manager being developed at General Atomics. We were pleased that COGENT was selected to represent an edge plasma code in this integration project. We expect our participation in AToM to help further ensure that the research results produced by this base program project will impact SC applications and programs.

References

- [1] <https://github.com/bendudson/Hypnotoad/>.
- [2] J. W. Banks, J. A. F. Hittinger, J. M. Connors, and C. S. Woodward. A posteriori error estimation via nonlinear error transport with application to shallow water. *Contemp. Math.*, 586:35–42, 2013.
- [3] R. H. Cohen, M. Dorf, and M. Dorr. Reduced electron models for edge simulation. *Contrib. Plasma Phys.*, 52(5–6):529–533, 2012.
- [4] P. Colella, M. R. Dorr, J. A. F. Hittinger, and D. F. Martin. High-order, finite-volume methods in mapped coordinates. *J. Comput. Phys.*, 230:2952–2976, 2011. Also available as Lawrence Livermore National Laboratory report LLNL-JRNL-385708.
- [5] C. G. Gibson. *Elementary Geometry of Differentiable Curves: An Undergraduate Introduction*. Cambridge University Press, New York, 2001.
- [6] P. McCorquodale, M. R. Dorr, J. A. F. Hittinger, and P. Colella. High-order finite-volume methods for hyperbolic conservation laws on mapped multiblock grids. *J. Comput. Phys.*, 288:181–195, 2015.